



Community performance optimization

Making your people run as smoothly as your site

Brion Vibber

CTO, Wikimedia Foundation (2005-2009)
Senior Software Architect, StatusNet

WikiSym - Orlando, FL

October 27, 2009

I) People are not computers


There was some concern when I came up with the title for my talk that people won't like being compared to software or server processes. I admit it sounds kind of impersonal, so I'd like to start with this disclaimer: people are not computers.

2) But computers are kinda like people


But... computers **are** kinda like people. People who are **very** good at doing what you tell them to, but **very** bad at doing what you meant them to do!

3) Thus we can learn some nice lessons


This is pretty important, because the power of people on any social website is circumscribed by the software that runs it. Web app developers need to understand how to make decisions that will benefit both the computers and their users... and there are a lot of problems that are similar for both.




Everybody talks these days about the end of Moore's Law and how it's a multicore world out there.




Even entry-level laptops are often dual-core today, and you can't make the most of your hardware unless the software splits up its jobs into smaller pieces which can be doled out to each processor.





Quad-core is beginning to take over the desktop, and a lot of developers are a little scared about where things are going in the future.



Of course those of us writing server-side software have been dealing with this problem for years. Web sites in particular are usually highly scalable, with front-end code that spreads the workload for hundreds, thousands, or millions of users over many CPUs on multiple machines.




24x web UI




1x database

We know well that to keep those systems running smoothly, we have to avoid bottlenecks where the parallelized processing depends on a limited resource.



24x web UI



4x database

If we have different classes of services that depend on each other, then we need to carefully balance them to ensure that the tubes stay clear.



This is just as true for social software -- maybe even more so. There's no Moore's law for people; we might have better tools but we're still just as human year after year.



But our projects and our sites get bigger...





...and bigger...



...and bigger.



Posters





Moderator

And we run into those same kinds of bottlenecks on forums and mailing lists...



Wiki contributors





Admin on
deletion patrol

...on wikis...






Patch contributors





SVN committer

...and in the human side of free software development.



Edit conflict!


Wikis were meant to counter exactly this sort of problem, by making it easy to for people to join in and act. In software terms, we can think of this as optimistic locking -- we only have to deal with a problem if it actually occurs.



“Hey page B is junk, not a real article.
Can we delete it?”

“No time to
look at it. :(”

But on massively-active sites like Wikipedia, we've seen that new social bottlenecks can develop in the resolution stage, as potentially irreversible admin actions limit the number of people who can act.



Page B
IS GAY



“AfD voted 100% in favor of deleting
page B. Can you delete it for us?”

“Sure thing!”

Wikipedia attempts to reduce the burden by running things like deletion requests through a discussion and voting queue, giving multiple editors the chance to look and comment on the validity of the request. Admins can then ignore the bogus requests and have more confidence in the legit requests.

Wikipedia:Articles for deletion

From Wikipedia, the free encyclopedia

[Administrator instructions](#)

Today's Article for deletion log



Before listing an article for deletion here, consider whether a more efficient alternative is appropriate:

- For [problems that do not require deletion](#), including duplicate articles, articles needing improvement, pages needing redirects, or POV problems, **be bold** and fix the problem or [tag the article](#) appropriately.
- If an article is a copyright violation, please list it at [Wikipedia:Copyright problems](#).
- Some articles may qualify for speedy deletion; please refer to the [speedy deletion criteria and process](#).
- For non-controversial deletions, please refer to the [proposed deletion process](#).
- For a potentially controversial merger, consider listing it at [proposed mergers](#).

[Purge the cache to refresh this page](#)

Articles for deletion (AfD) is where [Wikipedians](#) discuss whether or not an article should be deleted. Articles listed are debated for at least seven days, after which the [deletion process](#) proceeds based on community [consensus](#). Then the page may be kept, [merged](#) or [redirected](#), [transwikied](#) (copied to another [Wikimedia](#) project), [renamed/moved](#) to another title, [userfied](#) to a user subpage, or deleted per the [deletion policy](#).

Shortcut:
[WP:AFD](#)


[Deletion discussions](#)

[Deletion today](#)

Of course even then, the process can become a bottleneck as it grows. Whether articles should be deleted, when, by whom, and on what criteria remains an ugly battleground.

Wikipedia:Requests for deletion

From Wikipedia, the free encyclopedia



This page is currently inactive and is retained for historical reference. Either the page is no longer relevant, or consensus on its purpose has become unclear. To revive discussion regarding the subject, seek broader input via a forum such as the [village pump proposals page](#).

Article deletion on Wikipedia is discussed at [Wikipedia:Articles for deletion](#). See also [Wikipedia:Proposed deletion](#).

- *Centralised deletion reform discussion at [Wikipedia:Deletion reform](#)*

If you need a page deleted, add the **Category:Delete** tag to it, and present your reasons for wanting it removed on the article's discussion page. (See [Wikipedia:DPL deletion](#) for a proposed implementation of this.)

This policy is intended to replace [Wikipedia:Votes for deletion](#).

SPECIAL REQUEST: The term "Requests for deletion" is used by Wiktionary for its d was chosen to avoid any confusion with Wikipedia's process. If you are going to give the procedure please choose something else. [Eclecticology](#) 01:00:10, 2005-08-02 (UTC) (Wiktionary Admin)


But it's the best name :-). Why can't we both use it? [Pcb21l Pete](#) 10:39, 3 August 2005 (UTC)

Nomination for Deletion? (WP:NFD)? (not that I like the idea, but it's a better name than wikipedia acronyms I can think of) [Proto t c](#) 15:22, 3 August 2005 (UTC)

[Articles for Deletion \(WP:AFD\)](#), so as to parallel TFD, RFD, and CFD. ~~~~ [N](#) (t/c) 15:22, 3 August 2005 (UTC)

Wikipedia:Deletion reform

From Wikipedia, the free encyclopedia



This page is currently inactive and is retained for historical reference. Either the page is no longer relevant, or consensus on its purpose has become unclear. To revive discussion regarding the subject, seek broader input via a forum such as the [village pump proposals page](#).

The deletion procedure on Wikipedia, [Wikipedia:Articles for deletion](#), has recently come under harsh criticism from many users on Wikipedia, and there seems to be a general opinion in the community that it has to be somehow replaced or changed. This page is an effort to coordinate the different discussions and make it into a coherent, readable format.

Contents [\[hide\]](#)

- [1 Background](#)
- [2 Ideas/Suggestions](#)
- [3 Current proposals for a new system](#)
- [4 Historical proposals](#)
- [5 General comments](#)
- [6 See also](#)

Background

[\[edit\]](#)

There have been several attempts and proposals to revamp the system over the years, including at least two renames! But so far, we haven't made significant changes to it. The process is hard to find, and discussions are hard to join and laden with jargon and strange markup.

- Local CPU resources are limited
- Remote resources are **more** expensive
- Locking and protocol overhead kill scalability.


So what are the key principles to keep in mind for such reform? We know some of the keys to scalable technical systems...

- People have limited time & patience
- Waiting on other people to act is **slow**
- People want to do what interests them
...not deal with process!

These translate easily to humans!

- Get out of peoples' way and let them get stuff done!

Or to put it more simply...



We've done some experimentation with the development process on the MediaWiki engine that powers Wikipedia. We're using traditional SVN, but we've done two crazy things: we approve a relatively large number of committers for our project size (>100)...



site

- [Main Page](#)
- [Browse categories](#)
- [Community portal](#)
- [Recent changes](#)
- [Tech blog](#)

search

support

- [User help](#)
- [Technical manual](#)
- [FAQ](#)
- [Support desk](#)

download

- [Get MediaWiki](#)
- [Get extensions](#)

development

- [Bug tracker](#)
- [Download from SVN](#)
- [View code changes](#)
- [Code comments](#)
- [Browse SVN](#)
- [Code docs](#)

[special page](#)

Code Review

For [MediaWiki](#) ([recent comments](#) | [status changes](#) | [tags](#) | [authors](#) | [states](#) | [release notes](#))

Search revisions in this repo by path

Path:

Show items per page

▲Revision	Status	Notes	Tests	Path	Commit summary	Author	Date
57454	new	0	559/559	/trunk/phase3	* (bug 21030) Fixed schema choices from being overwritten by defining unique fie...	overlordq	08:47, 7 October 2009
57453	new	0	559/559	/trunk/phase3/includes/filerep...	Set default thumb cache to 24hrs, rather than disabled	demon	03:46, 7 October 2009
57441	ok	0	559/559	/trunk/phase3	Cleanup r57359 -- removing PHP4-isms (references not needed to pass objects unle...	brion	22:14, 6 October 2009
57437	deferred	0	559/559	/trunk/phase3	Localisation updates for core messages from translatewiki.net (2009-10-06 21:14 ...	siebrand	21:22, 6 October 2009
57434	ok	0	559/559	/trunk/phase3	* (bug 21026) Fixed file redirects on shared repos on non-English client wikis	brion	20:23, 6 October 2009
57433	deferred	0	559/559	/trunk/phase3/includes/Default...	typo fix in comment	brion	19:46, 6 October 2009
57429	ok	0	559/559	/trunk/phase3/includes/History...	Cleanup for r57288 : use \$wgContLang->alignEnd() instead of the isRtl() ? 'left' ...	brion	18:54, 6 October 2009
57421	ok	0	559/559	/trunk	Found some more source files.	tstarling	10:19, 6 October 2009
57420	ok	0	559/559	/trunk	(bug 19741) Moved the XCF files from phase3 to their own artwork directory, to r...	tstarling	08:40, 6 October 2009
					(bug 20966) Use the name of the		

...and we've built tools to aid in review of changes *post*-commit to keep development motion fast and agile. Trunk is slightly wobblier, but more people can push more good work straight into core.

				/messag...		
56459	ok	0	559/559	/trunk/phase3 /languages /classe...	Fix E_NOTICE in r55415 breaking Zh variants	werdna
56458	resolved	2	555/559	/trunk/phase3 /languages /classe...	Partial revert of r55415, calling wfMsg in Language object constructor causes un...	werdna
				/trunk/phase3		

Wikitext parser regression tests are run automatically on every commit to the MediaWiki core, alerting all developers. Since this has gone in, most such regressions are fixed very quickly -- previously they could persist for weeks before someone remembered to run the test suite.

Comments

#Comment by [PhiLiP](#) ([Talk](#) | [contribs](#)) 01:35, 17 September 2009

LanguageZh has constructed its parent object **before** the wfMsg was called. So I don't think the wfMsg called in LanguageZh object constructor would cause unstub loop. Otherwise, Language object has also provided functions that used wfMsg.

#Comment by [PhiLiP](#) ([Talk](#) | [contribs](#)) 01:49, 17 September 2009

Test show it should be my fault. I'll find another way to accomplish this. I'm sorry.

A built-in threaded discussion forum gives committers and non-committing devs alike the chance to discuss a change; in this case the author of a previous revision that interacted badly with this one explains how it went wrong. Many problems are found, hashed out, and fixed before the core committers get to their own code review!

Commit access requests

This page serves as a central location for users to request [commit access](#) to [SVN](#). Please note that just because you've listed your name here doesn't mean that you'll get access. Final approval for access still lies with [Brion](#) and [Tim](#).

Contents [\[hide\]](#)

- [1 Request format](#)
- [2 Completed requests](#)
- [3 Current requests](#)
 - [3.1 Tidan](#)
 - [3.2 InforMedic](#)
 - [3.3 Bile](#)
 - [3.4 Matěj Grabovský](#)
 - [3.5 User:Michael Daly2](#)
 - [3.6 Erich](#)
 - [3.7 MacMed](#)
 - [3.8 Marco 27](#)



We try to get MediaWiki developers on [board](#) before they grow up and get a real job.

Request format

[\[edit\]](#)


Copy and paste the following text to the bottom of the list. Of course you can add extra information you think is relevant, but try to keep it simple and to the point.

```
=== User name ===
* Requested commit name: <tt>abcd</tt>\[note 1\]
* Expected area to work in: extensions, localization, etc
* [http://example.com link to public key]\[note 2\]
* Past work on MediaWiki - patches, extensions, bug reports, etc\[note 3\]
```

It's still not perfect. Review of patches coming through our bug tracker, and commit access setup for extension authors & patchers have remained bottlenecked, and we've discussed clearing it up with a potential migration to git.

Classic CVS/SVN model


Patch submitters



The classic CVS or Subversion model for source control is inherently bottlenecked. A small group of core committers have direct access to the repository; any of them can commit directly, but all patches from other developers need to go through that core group for review.

Classic CVS/SVN model


Patch submitters



This makes it difficult for multiple people to be involved in a feature or bug fix during its development, since the change can only be shared by sending static patch files around. The burden of testing and review falls mainly on the few core committers.

Git model


Patch submitters



Used properly, distributed VCSes can reduce those bottlenecks. We've been using Git via Gitorious for StatusNet to good effect. With git, *everyone* gets the benefits of version control while they work, not just the core committers.



GITORIOUS

[All Projects](#)[All Teams](#)[~brion](#)  6[Logout](#) [StatusNet](#)[mainline](#)[Merge requests](#)

Merge requests in mainline

A "merge request" is a notification from one repository to another that would like their changes to be merged upstream.

Filter by status:

Open: [Open](#)

Closed: [Merged](#) [Rejected](#)

ID	Status	Summary	Target branch	Creator	Age
#1879	Merged	Force login to download file attachments for private sites	0.8.x	jeff-themovie	5 days
#1815	Merged	Adds support for design colors to the twitter API	0.9.x	tdfischer	12 days
#1738	Merged	Fixes related to sms/twitter enabled/disabled and openidonly	0.8.x	jeff-themovie	18 days
#1730	Merged	Don't show search suggestions for private sites	0.8.x	jeff-themovie	19 days
#1723	Merged	Fixed various E_NOTICES	0.8.x	jeff-themovie	20 days
#1633	Merged	fix for a misnamed variable in subscription function in the "create simulation data" script	0.9.x	julien	about 1 month
#1631	Merged	Fix URL linkification test cases for addition of 'title' attribution with long URL in f3c8fccc	0.9.x	brion	about 1 month
#1569	Merged	Fixes for posting shortened URLs or uploads	0.9.x	brion	about 1 month
#1568	Merged	Don't trigger E_NOTICE when looking for commands in the notice input	0.9.x	brion	about 1 month
#1525	Merged	Pagination of user/tag went to wrong page.	0.8.x	cvollick	about 1 month

Ideally, working on code through your Git clone is easier than running patches through a bug tracker. It's easier to share your work with other coders and testers, and tools like Gitorious and Github make the final review-and-merge one-click easy for core committers.

All filters

Options

Deleted filters: ☐ Include deleted filters ☒ Hide deleted filters ☐ Show only deleted filters

Disabled filters: ☐ Hide disabled filters

Number per page: 50

Update



First page



Previous page



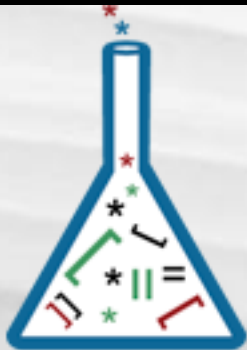
Next page



Last page

▼ Filter ID	Public description	Consequences	Status	Last modified	Visibility	Hit count
1	Prodego's test filter		Disabled	00:15, 21 September 2009 by Prodego (talk contribs)	Public	867 hits
3	New user blanking articles	Warn, Tag	Enabled	20:44, 15 October 2009 by Rich Farmbrough (talk contribs)	Public	75,590 hits
5	User self-renaming or moving user talk pages into article talk space	Warn, Tag	Enabled	08:47, 20 June 2009 by Ruslik0 (talk contribs)	Public	497 hits
9	Personal attacks by unregistered or new user	Warn, Disallow	Enabled	15:06, 29 September 2009 by Ruslik0 (talk contribs)	Public	26,027 hits
10	Unusual change in title length	Warn, Disallow	Enabled	18:10, 29 August 2009 by NawlinWiki (talk contribs)	Private	166 hits
11	You/He/She/It sucks	Warn, Tag	Enabled	12:00, 3 October 2009 by Ruslik0 (talk contribs)	Public	11,283 hits
12	Replacing a page with obscenities	Warn, Disallow	Enabled	02:56, 23 October 2009 by Prodego (talk contribs)	Private	13,818 hits
17	Prolific socker	Disallow	Enabled	11:10, 26 October 2009 by Wknight94 (talk contribs)	Private	280 hits
18	Test type edits from clicking on edit bar	Warn, Tag	Enabled	05:38, 24 July 2009 by King of Hearts (talk contribs)	Public	114,297 hits

So how can we translate these experiments back to wikiland? One of the first major things we pushed to the insanely-active English-language Wikipedia this year was an automated filtering system which can provide pre-edit feedback or outright prevent the most frequent types of problem edits, giving our pool of administrators more time to spend on real issues.



WIKIMEDIA
LABORATORIES

navigation

- [Main Page](#)
- [FlaggedRevs issues](#)
- [Recent changes](#)
- [Random page](#)

search

[Go](#) [Search](#)

toolbox

- [What links here](#)
- [Related changes](#)
- [Special pages](#)
- [Printable version](#)
- [Permanent link](#)

in other languages

- [Deutsch](#)
- [Español](#)
- [Français](#)
- [Simple English](#)
- [中文](#)


Help us improve Wikipedia by [supporting it financially](#).

[page](#) [draft](#) [discussion](#) [edit draft](#) [history](#)

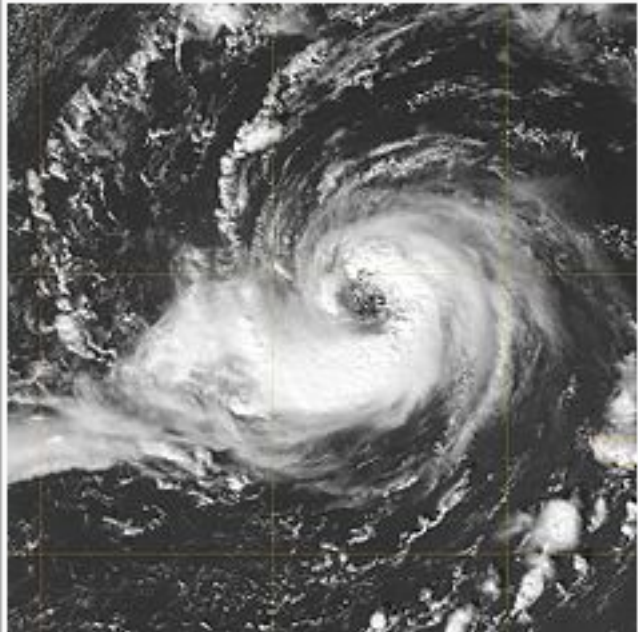
Welcome to the **flagged protection** test! • Please report any issues / bugs to [FlaggedRevs issues](#).

Hurricane Vince (2005)

From Wikimedia

 This is the latest checked version (newest draft) (+/-)

Hurricane Vince
Category 1 hurricane (**SSHS**)



Hurricane Vince on October 9 14:23 UTC

Formed	October 8, 2005
Dissipated	October 11, 2005
Highest winds	120 km/h (75 mph) (1-minute sustained)
Lowest pressure	988 hPa (mbar)
Fatalities	None reported
Damage	Unknown
Areas	Madeira Islands , southern

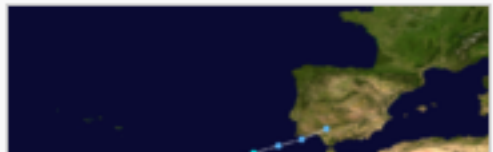
Hurricane Vince was an unusual **hurricane** which developed in the northeastern **Atlantic basin**. Forming in October during the **2005 Atlantic hurricane season**, the waters over which it developed were considered too cold for tropical development. Vince was the 20th **named tropical cyclone** and 12th **hurricane** of the extremely active season.^[*citation needed*]

Vince developed from a non-tropical system on October 8, becoming a **subtropical storm** southeast of the **Azores**. The **National Hurricane Center** (NHC) did not officially name the storm until the next day, shortly before Vince became a hurricane. The storm weakened at sea and, on October 11, made landfall on the **Iberian Peninsula** as a tropical depression. Vince was the first tropical system to do so since the **1842 Spain hurricane**. It dissipated over Spain, bringing much needed rain to the region, and its remnants passed into the **Mediterranean Sea**.

Contents [hide]

- 1 [Meteorological history](#)
- 2 [Preparations and impact](#)
- 3 [Records and naming](#)
- 4 [See also](#)
- 5 [References](#)
- 6 [External links](#)

Meteorological history [edit]



On October 500, an operationally **unnamed subtropical storm** which had gone unnoticed by the NHC was absorbed by a non-tropical frontal **low**, which was moving to the southeast over the

We're also working on variants of the Flagged Revisions system, introduced in 2007 to the German-language Wikipedia and used on a number of other wikis. For English Wikipedia we're testing a configuration that interferes as little as possible in the general workflow, concentrating on replacing page protection for articles that are frequently mis-edited.

Editing Hurricane Vince (2005)

From Wikimedia

This page is subject to [flagged protection](#).

- 19:51, 9 October 2009 [Mentifisto](#) ([Talk](#) | [contribs](#)) configured stable versioning for [Hurricane Vince \(2005\)](#) (*[Precedence: None, Default: Stable](#)*)

[View full log](#)



Latest checked revision is at *20 October 2009*. [2 newer edits](#) need checking.

([past reviewed versions](#) • [review log](#))

[Confirmed version]

==Meteorological history==

{{storm path|Vince 2005 track.png}}

On October **500**, an operationally [\[\[2005 Azores subtropical storm|unnamed subtropical storm\]\]](#) which had gone unnoticed by the NHC was absorbed by a non-tropical frontal [\[\[Low pressure area|low\]\]](#), which was moving to the southeast over the Azores. The low pressure system gained a more concentrated circulation and lost its [\[\[Weather front|frontal structure\]\]](#) after absorbing the subtropical storm. The

[Draft version]

==Meteorological history==

{{storm path|Vince 2005 track.png}}

On October **2005**, an operationally [\[\[2005 Azores subtropical storm|unnamed subtropical storm\]\]](#) which had gone unnoticed by the NHC was absorbed by a non-tropical frontal [\[\[Low pressure area|low\]\]](#), which was moving to the southeast over the Azores. The low pressure system gained a more concentrated circulation and lost its [\[\[Weather front|frontal structure\]\]](#) after absorbing the subtropical storm. The

Flagged-protected pages may be freely edited, with only the approval for general display of the updated draft needing administrator intervention post-edit -- a much smaller burden than requiring admins to make edits on other users' behalf!

Extension:DeleteQueue

What can this extension do?

[\[edit\]](#)

[Manual on MediaWiki Extensions](#)
[List of MediaWiki Extensions](#)

Creates a queue-based system for managing deletion.

Installation

[\[edit\]](#)

To install this extension, add the following to [LocalSettings.php](#):

```
#add configuration parameters here  
#setup user rights here  
require_once( "$IP/extensions/DeleteQueue/DeleteQueue.php" );
```



DeleteQueue

Release status: unknown

Description Creates a queue-based system for managing deletion

Author(s) Andrew Garrett ([Werdna^{Talk}](#))

License GPL 2 or later

Download [Download snapshot](#)
[Subversion](#) [\[Help\]](#)
[Browse source code](#)

[check usage](#) [\[Help\]](#) *(experimental)*

The deletion discussion problem I showed earlier is well-known, and unfortunately we haven't fixed it yet. There has been some work on an extension to provide a dedicated communications channel for it, which should make it easier for people to find out what's going on with their article and be involved with what happens to it once it's ready.

Where is the work from? (Click on the appropriate link)

- It is [entirely my own work](#)
- It is [someone else's work from Flickr](#) (more information on uploading Flickr images)
- It is [a derivative work of a file from Commons](#)
- It is [from another Wikimedia project](#) (Wikipedia, Wikibooks, Wikinews, etc.)
- It is [from a US federal government source](#)
- It is [a free software screenshot](#)
- It is [from somewhere else](#)
- It is [a non-free fair use image](#) (Non-free logo, cover, screenshot, artwork etc.)
- I don't know who the author is, or I don't know what license applies
- I need help figuring out what the license is



For experienced users:

- Already know the license, and its copyright tag? Go directly to the [main upload form](#).
- Already have an `{{information}}` template ready? (For instance, generated by the [Finfo](#) tool?) Use the [basic upload form](#).

See also the full list of [copyright tags](#) for many nations. The [first section](#) of that page also has other general tags for public domain sources.


The process for adding image and media files is perhaps one of the worst currently. Not only do we have all the same problems with review for copyright and notability after your upload, but you have to dig through all this first...

Upload your own work

This form is only for uploading your *own work* - files that **you have created** and where **you own the copyright**.

To upload work created by someone else, please use **a different form**.



Step 1. Is this your work?

- Photographs or videos **you** have created of:
 - natural landscapes, animals, and plants
 - people that are either public figures or are taken at public events
 - useful or non-artistic objects (tools, dinner plates, etc.)
-  Graphs, maps, diagrams, and audio you have created entirely yourself.
- Works that are under a **free license** that **you** have modified:
 - like above but the original work was not yours
 - the original author(s) need to be credited according to the license terms

- Works which contain other copyrighted works (**derivative works**) - these are **not permitted**, so don't upload them!
 - Screenshots of TV shows, movies, DVDs, and software.
 - Photographs of art, statues, commercial packaging and often toys.
 - Drawings of copyrighted characters from TV, Comics, or Movies, *even if you drew them yourself*.



Step 2. Help people find this file. *(Commons has over five million files)*

- Use a **descriptive destination filename**. *Don't* use default filenames generated by cameras!
- Write a good description of the contents of the file. Search engines search text, not image pixels!
- Add categories to the image. Use the [CommonSense](#)  tool to find **categories**.
- Consider adding location data to your image, see [Commons:Geocoding](#) for more information.
- If you already have an `{{information}}` template or similar template ready to copy/paste, use the **basic upload form** . To *always* use the old basic form layout, check "Use the old-style upload form layout." on the "Gadgets" tab of *your preferences*.



Even the directions take up two pages...

When you upload your work to commons you are donating it to the world by using a free content license which allow everyone to **use**, **modify**, and **redistribute** your work for any purpose. **This donation is non-revocable**. Although you retain the copyright to your work, all works submitted to commons must have [at least these freedoms](#). To learn more about the charitable mission of the Wikimedia Foundation visit [the foundation website](#).



- Commons offers multiple licensing options. They differ primarily in their attribution requirements and the existence of additional terms that require works based on your work to be under similar free terms.
- Choose a license from the drop-down box, or write one in the summary box (from [Commons:Copyright tags](#)).
- It is your responsibility to understand what rights you are keeping for yourself.
 - More information can be found at [Commons:Licensing](#).
 - If you need help understanding licensing terms, please ask at the [Commons:Help desk](#).

Step 4. Has your work been published before? *(if so, help us verify that this is a rightful copy)*


- If your work is as yet unpublished, you can proceed with the upload.
- If your work was published on the internet under a free license and a credit to the same name or nickname as [you are using now](#), you can proceed with the upload.
- If your file was published without a free license, or with a different credit name, or with no credit at all, please send a confirmation E-mail to the address on [this page](#) and add `{{OTRS pending|month=October|day=27|year=2009}}` with details about the previous publication in the permission field below.

Other tips

- Upload the highest resolution file you have, and **don't watermark** your work.
- Only upload file types which [Commons accepts](#) (*SVG, PNG, JPG, GIF, Ogg, DjVU, ...*).
- Compromising or embarrassing images of non-public people taken without their knowledge are often problematic. Use good judgment.
- Wikimedia Commons is for [educational or informational content](#).


... did you read all that? I know I didn't.

Upload your own work

Local filename: 


Maximum file size: 100 MB


Permitted file types: png, gif, jpg, jpeg, xcf, mid, ogg, ogv, svg, djvu, tif, tiff, oga.


Destination filename: 


Name of the file at Commons after the upload. Use a descriptive file name ([tips](#)).

Original source: 


 Where does this file come from?


Author(s): 


 Who created this file? Also, if it shows some artwork, who created that?



Date of the work: 

Date of creation and/or first publication of the work ("YYYY-MM-DD" format).

Description: 

English 

Other versions: 

Permission:  Is this work of yours already published elsewhere? Use `{{OTRS pending|month=October|day=27|year=2009}}` and send permission by [e-mail](#) . Also for specialized [license tags](#).

If even longtime Wikipedians like myself can't be bothered to figure this form out, we've got a serious problem.

People want to get from point A to point B...

...remove the roadblocks!

<http://leuksman.com/pages/presentations>